

中华人民共和国国家知识产权局
STATE INTELLECTUAL PROPERTY OFFICE
OF THE PEOPLE'S REPUBLIC OF CHINA



证 明
CERTIFICATE

本证明之附件是向中国专利局作为受理局提交的下列国际申请副本
IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY OF THE BELOW
IDENTIFIED INTERNATIONAL APPLICATION THAT WAS FILED WITH THE
CHINESE PATENT OFFICE AS RECEIVING OFFICE

申请号: PCT/CN2006/000939

INTERNATIONAL APPLICATION NUMBER

申请日: 10. MAY 2006 (10. 05. 2006)

INTERNATIONAL FILING DATE

名称: THREAD SCHEDULING ON MULTIPROCESSOR SYSTEMS

INVENTION

**CERTIFIED COPY OF
PRIORITY DOCUMENT**

中华人民共和国国家知识产权局局长
COMMISSIONER OF THE STATE INTELLECTUAL PROPERTY
OFFICE OF THE PEOPLE'S REPUBLIC OF CHINA

二零零六年六月九日

JUNE 9, 2006

BEST AVAILABLE COPY

HOME COPY

PCT

REQUEST

The undersigned requests that the present international application be processed according to the Patent Cooperation Treaty.

For receiving Office use only

PCT/CN 2006 / 0 0 0 9 3 9
International Application No.

10 · MAY 2006 (10 · 05 · 2006)
International Filing Date

RO/CN 中华人民共和国国家知识产权局
PCT International Application
Name of receiving Office and "PCT International Application"

Applicant's or agent's file reference
(if desired) (12 characters maximum) FPEL06150029

Box No. I TITLE OF INVENTION	
THREAD SCHEDULING ON MULTIPROCESSOR SYSTEMS	
Box No. II APPLICANT <input type="checkbox"/> This person is also inventor	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.) INTEL CORPORATION 2200 Mission College Blvd. Santa Clara, California 95052 United States of America	Telephone No. Facsimile No. Teleprinter No. Applicant's registration No. with the Office
State (that is, country) of nationality: US	State (that is, country) of residence: US
This person is applicant for the purposes of: <input type="checkbox"/> all designated States <input checked="" type="checkbox"/> all designated States except the United States of America <input type="checkbox"/> the United States of America only <input type="checkbox"/> the States indicated in the Supplemental Box	
Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.) LI, Wenlong Room 18-4-502, Tiantongyuan West 2 District, Changping District, Beijing 100043 P. R. of China	This person is: <input type="checkbox"/> applicant only <input checked="" type="checkbox"/> applicant and inventor <input type="checkbox"/> inventor only (If this check-box is marked, do not fill in below.) Applicant's registration No. with the Office
State (that is, country) of nationality: CN	State (that is, country) of residence: CN
This person is applicant for the purposes of: <input type="checkbox"/> all designated States <input type="checkbox"/> all designated States except the United States of America <input checked="" type="checkbox"/> the United States of America only <input type="checkbox"/> the States indicated in the Supplemental Box	
<input checked="" type="checkbox"/> Further applicants and/or (further) inventors are indicated on a continuation sheet.	
Box No. IV AGENT OR COMMON REPRESENTATIVE; OR ADDRESS FOR CORRESPONDENCE	
The person identified below is hereby/has been appointed to act on behalf of the applicant(s) before the competent International Authorities as: <input type="checkbox"/> agent <input type="checkbox"/> common representative	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country.) China Patent Agent (H.K.) Ltd. 22/F, Great Eagle Centre 23 Harbour Road, Wanchai Hong Kong Special Administrative Region The People's Republic of China	Telephone No. (852)28284688 Facsimile No. (852)28271018 Teleprinter No. Agent's registration No. with the Office
<input type="checkbox"/> Address for correspondence: Mark this check-box where no agent or common representative is/has been appointed and the space above is used instead to indicate a special address to which correspondence should be sent.	

Continuation of Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)

If none of the following sub-boxes is used, this sheet should not be included in the request.

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

LI, Eric Q.
Room 3702, #4 North of Xibahe,
Chaoyang District, Beijing 100028
P. R. of China

This person is:

- ☐ applicant only
- ☒ applicant and inventor
- ☐ inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

CN

State (that is, country) of residence:

CN

This person is applicant for the purposes of:

☐ all designated States☐ all designated States except the United States of America☒ the United States of America only☐ the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

DI, Nan
Peking Univerisity, Haidian District
Beijing 100000
P. R. of China

This person is:

- ☐ applicant only
- ☒ applicant and inventor
- ☐ inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

CN

State (that is, country) of residence:

CN

This person is applicant for the purposes of:

☐ all designated States☐ all designated States except the United States of America☒ the United States of America only☐ the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

This person is:

- ☐ applicant only
- ☐ applicant and inventor
- ☐ inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

State (that is, country) of residence:

This person is applicant for the purposes of:

☐ all designated States☐ all designated States except the United States of America☐ the United States of America only☐ the States indicated in the Supplemental Box

Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.)

This person is:

- ☐ applicant only
- ☐ applicant and inventor
- ☐ inventor only (If this check-box is marked, do not fill in below.)

Applicant's registration No. with the Office

State (that is, country) of nationality:

State (that is, country) of residence:

This person is applicant for the purposes of:

☐ all designated States☐ all designated States except the United States of America☐ the United States of America only☐ the States indicated in the Supplemental Box

Further applicants and/or (further) inventors are indicated on another continuation sheet.

Box No. V DESIGNATIONS

The filing of this request constitutes under Rule 4.9(a), the designation of all Contracting States bound by the PCT on the international filing date, for the grant of every kind of protection available and, where applicable, for the grant of both regional and national patents. However,

- ☐ DE Germany is not designated for any kind of national protection
☐ JP Japan is not designated for any kind of national protection
☐ KR Republic of Korea is not designated for any kind of national protection
☐ RU Russian Federation is not designated for any kind of national protection

(The check-boxes above may only be used to exclude (irrevocably) the designations concerned if, at the time of filing, the international application contains in Box No. VI a priority claim to an earlier national application filed in the particular State concerned, in order to avoid the ceasing of the effect, under the national law, of this earlier national application. See the Notes to Box No. V as to the consequences of such national law provisions in these States).

Box No. VI PRIORITY CLAIM

The priority of the following earlier application(s) is hereby claimed:

Filing date of earlier application (day/month/year)	Number of earlier application	Where earlier application is:		
		national application: country or Member of WTO	regional application:* regional Office	international application: receiving Office
item (1)				
item (2)				
item (3)				

☐ Further priority claims are indicated in the Supplemental Box.

The receiving Office is requested to prepare and transmit to the International Bureau a certified copy of the earlier application(s) *(only if the earlier application was filed with the Office which for the purposes of this international application is the receiving Office)* identified above as:

☐ all items ☐ item (1) ☐ item (2) ☐ item (3) ☐ other, see Supplemental Box

** Where the earlier application is an ARIPO application, indicate at least one country party to the Paris Convention for the Protection of Industrial Property or one Member of the World Trade Organization for which that earlier application was filed (Rule 4.10(b)(ii)):*

Box No. VII INTERNATIONAL SEARCHING AUTHORITY

Choice of International Searching Authority (ISA) *(if two or more International Searching Authorities are competent to carry out the international search, indicate the Authority chosen; the two-letter code may be used):*

ISA /

Request to use results of earlier search; reference to that search *(if an earlier search has been carried out by or requested from the International Searching Authority):*

Date (day/month/year) Number Country (or regional Office)

Box No. VIII DECLARATIONS

The following declarations are contained in Boxes Nos. VIII (i) to (v) *(mark the applicable check-boxes below and indicate in the right column the number of each type of declaration):*

Number of
declarations

- | | | |
|---|--|---|
| <input type="checkbox"/> Box No. VIII (i) | Declaration as to the identity of the inventor | : |
| <input type="checkbox"/> Box No. VIII (ii) | Declaration as to the applicant's entitlement, as at the international filing date, to apply for and be granted a patent | : |
| <input type="checkbox"/> Box No. VIII (iii) | Declaration as to the applicant's entitlement, as at the international filing date, to claim the priority of the earlier application | : |
| <input type="checkbox"/> Box No. VIII (iv) | Declaration of inventorship (only for the purposes of the designation of the United States of America) | : |
| <input type="checkbox"/> Box No. VIII (v) | Declaration as to non-prejudicial disclosures or exceptions to lack of novelty | : |

Box No. IX CHECK LIST; LANGUAGE OF FILING

This international application contains:		This international application is accompanied by the following item(s) (mark the applicable check-boxes below and indicate in right column the number of each item):	Number of items
(a) on paper, the following number of sheets:		1. <input checked="" type="checkbox"/> fee calculation sheet	: 1
request (including declaration sheets)	: 4	2. <input checked="" type="checkbox"/> original separate power of attorney	: 1
description (excluding sequence listing and/or tables related thereto)	: 19	3. <input type="checkbox"/> original general power of attorney	:
claims	: 7	4. <input type="checkbox"/> copy of general power of attorney; reference number, if any:	:
abstract	: 1	5. <input type="checkbox"/> statement explaining lack of signature	:
drawings	: 8	6. <input type="checkbox"/> priority document(s) identified in Box No. VI as item(s):	:
Sub-total number of sheets	: 39	7. <input type="checkbox"/> translation of international application into (language):	:
sequence listing	:	8. <input type="checkbox"/> separate indications concerning deposited microorganism or other biological material	:
tables related thereto	:	9. <input type="checkbox"/> sequence listing in electronic form (indicate type and number of carriers)	:
(for both, actual number of sheets if filed on paper, whether or not also filed in electronic form; see (c) below)	:	(i) <input type="checkbox"/> copy submitted for the purposes of international search under Rule 13ter only (and not as part of the international application):	:
Total number of sheets	: 39	(ii) <input type="checkbox"/> (only where check-box (b)(i) or (c)(i) is marked in left column) additional copies including, where applicable, the copy for the purposes of international search under Rule 13ter	:
		(iii) <input type="checkbox"/> together with relevant statement as to the identity of the copy or copies with the sequence listing mentioned in left column	:
(b) <input type="checkbox"/> only in electronic form (Section 801(a)(i))		10. <input type="checkbox"/> tables in electronic form related to sequence listing (indicate type and number of carriers)	:
(i) <input type="checkbox"/> sequence listing		(i) <input type="checkbox"/> copy submitted for the purposes of international search under Section 802(b-quater) only (and not as part of the international application)	:
(ii) <input type="checkbox"/> tables related thereto		(ii) <input type="checkbox"/> (only where check-box (b)(ii) or (c)(ii) is marked in left column) additional copies including, where applicable, the copy for the purposes of international search under Section 802(b-quater)	:
(c) <input type="checkbox"/> also in electronic form (Section 801(a)(ii))		(iii) <input type="checkbox"/> together with relevant statement as to the identity of the copy or copies with the tables mentioned in left column	:
(i) <input type="checkbox"/> sequence listing		11. <input type="checkbox"/> other (specify):	:
(ii) <input type="checkbox"/> tables related thereto			
Type and number of carriers (diskette, CD-ROM, CD-R or other) on which are contained the			
<input type="checkbox"/> sequence listing:			
<input type="checkbox"/> tables related thereto:			
(additional copies to be indicated under items 9(ii) and/or 10(ii), in right column)			
Figure of the drawings which should accompany the abstract:		Language of filing of the international application:	EN

Box No. X SIGNATURE OF APPLICANT, AGENT OR COMMON REPRESENTATIVE

Next to each signature, indicate the name of the person signing and the capacity in which the person signs (if such capacity is not obvious from reading the request).



For receiving Office use only		2. Drawings: <input type="checkbox"/> received: <input type="checkbox"/> not received:
1. Date of actual receipt of the purported international application:	10 · MAY 2006 (10 · 05 · 2006)	
3. Corrected date of actual receipt due to later but timely received papers or drawings completing the purported international application:		
4. Date of timely receipt of the required corrections under PCT Article 11(2):		
5. International Searching Authority (if two or more are competent): ISA /	6. <input type="checkbox"/> Transmittal of search copy delayed until search fee is paid	

For International Bureau use only
Date of receipt of the record copy by the International Bureau:

PCT

FEE CALCULATION SHEET

Annex to the Request

For receiving Office use only

PCT/CN 2006 / 0 0 0 9 3 9
International Application No.

10 · MAY 2006 (10 · 0 · 5 · 2006)
Date stamp of the receiving Office

Applicant's or agent's
file reference

FPEL06150029

Applicant

INTEL CORPORATION

CALCULATION OF PRESCRIBED FEES

1. TRANSMITTAL FEE

CNY500

T

CNY 500. ✓

2. SEARCH FEE

CNY2100

S

CNY 2100. ✓

International search to be carried out by

CN

(If two or more International Searching Authorities are competent to carry out the international search, indicate the name of the Authority which is chosen to carry out the international search.)

3. INTERNATIONAL FILING FEE

Where items (b) and/or (c) of Box No. IX apply, enter Sub-total number of sheets } 39
Where items (b) and (c) of Box No. IX do not apply, enter Total number of sheets }

i1 first 30 sheets

CHF1400

i1

CHF 1400. ✓

i2 9
number of sheets
in excess of 30

x

CHF15
fee per sheet

=

CHF135

i2

CHF-135. ✓

i3 additional component (only if sequence listing and/or tables related thereto are filed in computer readable form under Section 801(a)(i), or both in that form and on paper, under Section 801(a)(ii)):

400 x
fee per sheet

i3

Add amounts entered at i1, i2 and i3 and enter total at I

CHF1535

I

CHF-1535. ✓

(Applicants from certain States are entitled to a reduction of 75% of the international filing fee. Where the applicant is (or all applicants are) so entitled, the total to be entered at I is 25% of the international filing fee.)

4. FEE FOR PRIORITY DOCUMENT (if applicable)

P

CNY 2600. ✓

5. TOTAL FEES PAYABLE

Add amounts entered at T, S, I and P, and enter total in the TOTAL box

CNY2600CHF1535

TOTAL

CHF-1535. ✓

MODE OF PAYMENT

☒ authorization to charge
deposit account (see below)

☐ postal money order

☒ cash

☐ coupons

☐ cheque

☐ bank draft

☐ revenue stamps

☐ other (specify):

AUTHORIZATION TO CHARGE (OR CREDIT) DEPOSIT ACCOUNT (This mode of payment may not be available at all receiving Offices)

☒ Authorization to charge the total fees indicated above.

☒ (This check-box may be marked only if the conditions for deposit accounts of the receiving Office so permit) Authorization to charge any deficiency or credit any overpayment in the total fees indicated above.

☒ Authorization to charge the fee for priority document.


Receiving Office: RO/ CN

Deposit Account

Date:

Name:

Signature:



THREAD SCHEDULING ON MULTIPROCESSOR SYSTEMS

BACKGROUND

1. FIELD

This disclosure relates generally to multi-thread applications on multi-processor systems, and, more specifically but not exclusively, to thread scheduling on a multi-processor system.

2. DESCRIPTION

A threaded application usually has data shared among its threads when running on symmetric multiprocessors ("SMP") and/or chip multiprocessors ("CMP"). The data sharing among different threads may be achieved in different ways but frequently done through a shared system-level memory. In a typical memory hierarchy in a multiprocessor system, a shared system-level memory between different processing cores has longer access latency for a processing core than a local cache of the processing core. Additionally, traffic (including coherency traffic) among different processing cores generated by excessive access to a shared system-level memory may saturate the bandwidth of the system interconnect (e.g., bus, ring, mesh, etc.). Therefore, it is desirable to investigate the data sharing behavior among different threads and reduce the cost of data transfer among threads.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the disclosed subject matter will become apparent from the following detailed description of the subject matter in which:

Figure 1 is a block diagram of an example multiprocessor system that uses a data sharing aware thread scheduling module, according to the disclosed subject matter in the present application;

Figure 2 is a block diagram of another example multiprocessor system that uses a data sharing aware thread scheduling module, according to the disclosed subject matter in the present application;

Figure 3 shows a block diagram of a data sharing aware thread scheduling module, according to the disclosed subject matter in the present application;

Figure 4 is a flowchart illustrating an example process for scheduling threads to target processors using information of data sharing behavior among different threads, according to the disclosed subject matter in the present application;

Figure 5 illustrates an example program that can be multithreaded in a multiprocessor system;

Figure 6 illustrates a code corresponding to the program shown in Figure 5, which is multithreaded for a multiprocessor system;

Figure 7 illustrates an example assembly code of the multi-threaded program shown in Figure 6;

Figure 8 illustrates an example assembly code of the multi-threaded program shown in Figure 6, which uses a data sharing aware thread scheduling method, according to the disclosed subject matter in the present application;

Figure 9 illustrates an example code of a multi-threaded video mining program, according to the disclosed subject matter in the present application; and

Figure 10 illustrates performance improvement of a multithread application running on a multiprocessor system, obtained by using a data sharing aware thread scheduling module, according to the disclosed subject matter in the present application.

DETAILED DESCRIPTION

According to embodiments of the subject matter disclosed in this application, a compiler in a multiprocessor system may compile a received multithreaded application, analyze data sharing behavior among multiple threads of the multithreaded application, and provide such information to a thread scheduler in the multiprocessor system. At run time, the thread scheduler may group together threads that share data frequently based on the data sharing information provided by the compiler, and schedule threads in the same group to processors in the same cluster. Processors in the same cluster have a shared storage device and have shorter access latency to the shared storage device than to the system-level memory. If there are not enough available processors in the same cluster, the rest of the threads may

be assigned to processors that are electronically in proximity to the cluster. Additionally, a feedback module may collect information on data sharing behavior among threads during run time and feedback such information to the thread scheduler. The thread scheduler may use the feedback information to regroup and reschedule the threads to processors at the next available scheduling time.

Reference in the specification to “one embodiment” or “an embodiment” of the disclosed subject matter means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosed subject matter. Thus, the appearances of the phrase “in one embodiment” appearing in various places throughout the specification are not necessarily all referring to the same embodiment.

Figure 1 is a block diagram of an example multiprocessor system 100 that uses a data sharing aware thread scheduling module. System 100 may comprise multiple processors such as processor A (132A). Processors in system 100 may be connected to each other using a system interconnect 110. System interconnect 110 may be a Front Side Bus (FSB). Each processor may be connected to Input/Output (IO) devices as well as system-level memory 120 through the system interconnect. The system-level memory may be a Dynamic Random Access Memory (DRAM), a Synchronous DRAM (SDRAM), a Double Data Rate (DDR) SDRAM, or other types of memory devices. Each processor (e.g., processor A (132A)) may have its own cache with one or more levels (e.g., level 1 (L1) and level 2 (L2) caches, not shown

in the figure). In one embodiment, several processors (e.g., processor A (132A) through processor M (132M)) may have a shared local-level storage device such as cache 138. Cache 138 may be connected with processor 132A through processor 132M via a connection mechanism 136. Connection mechanism 136 may be a local interconnect (e.g., local bus), a connection ring, or a cross-bar.

Processors 132A through 132M, connection mechanism 136, and cache 138, together may form a cluster (cluster 1 (130)). Cluster 1 may be connected with other clusters (e.g., cluster N (140)), individual processors (not shown in the figure), IO devices and system-level memory 130, and other devices via system interconnect 110. In one embodiment, an individual processor (not shown in the figure) may have multiple processing cores. Each processing core may have its own cache and all of the processing cores together may have shared local cache. Such a processor with multiple processing cores may be treated similarly as a cluster (e.g., cluster 1). For the convenience of description, a processing core inside a processor may be treated in the same way as a single-core processor. Thus, the word "processor" may be used to represent either a single-core processor or a processing core in a multi-core processor.

A data sharing aware thread scheduling module (e.g., 134A, 134M) may be included in a processor (such as processor 132A or 132M). The thread scheduling module may comprise a compiler, a thread scheduler, and a feedback module. The compiler may receive a multithreaded application and compile the application into one or more object code which is specific to

the underlying architecture. The thread scheduler may be a part of, or associate with, an operating system (OS). The thread scheduler schedules multiple threads of the multithreaded application to different processors. The feedback module may observe actual data sharing behavior among threads during execution and provide the actual data sharing information to the thread scheduler, which will use the feedback information to fine tune its thread scheduling next time.

When two or more threads share data frequently, such data sharing is often achieved through shared system-level memory (e.g., memory 120). Access to a system-level memory has longer latency by a processor (e.g., processor 132A) than access to a local-level memory (e.g., cache 138) by the processor. Additionally, frequent access to a system-level memory consumes quite a bit bandwidth of the system interconnect (e.g., 110) and may even cause congestion on the system interconnect. Furthermore, coherence traffic generated by access to a system-level memory typically has high overhead.

A conventional thread scheduler does not have much knowledge on data sharing behavior among multiple threads, and when it schedules multiple threads, it mainly focuses on memory hierarchy optimization. According to one embodiment of the subject matter in the present application, the thread scheduler may utilize information on data sharing behavior among threads provided by the compiler, which is able to identify such information when compiling a multithreaded application. Additionally, the thread scheduler has the ability of knowing the underlying architecture of the system, for example, which processors are in the same cluster and which clusters are in proximity

to each other electronically. Clusters that are in proximity to each other electronically include such clusters where data sharing among their processors requires less time, causes less traffic on the system interconnect, and/or incurs less coherence overhead, than data sharing among processors from other clusters. Such clusters may include those that are close to each other in topology. With data sharing information from the compiler plus knowledge of underlying architectural characteristics, the thread scheduler (e.g., 134A) may schedule those tightly coupled threads (with frequent data sharing among them) to processors in the same cluster. If there are not enough available processors in any cluster, the thread scheduler may schedule those tightly coupled threads to processors in those clusters that are in proximity to each other electronically.

Figure 2 is a block diagram of another example multiprocessor system 200 that uses a data sharing aware thread scheduling module. In system 200, system interconnect 210 that connects multiple clusters (e.g., 220A, 220B, 220C, and 220D) is a links-based point-to-point connection. Each cluster may connect to the system interconnect through a links hub (e.g., 230A, 230B, 230C, and 230D). In some embodiments, a links hub may be co-located with a memory controller, which coordinates traffic to/from a system memory (not shown in the figure). Similar to cluster 1 (130) shown in Figure 1, clusters 220A, 220B, 220C, and 220D may comprise two or more processors (or processing cores), which are connected with a shared local storage device (e.g., last level cache). Each processor may have its own local caches.

According to one embodiment of the subject matter disclosed in the present application, the compiler may provide the thread scheduler data sharing information among threads. The thread scheduler may also obtain the underlying hardware configuration from the OS or a hardware abstract layer (e.g., virtual machine). Based on the data sharing information and the hardware structural information, the thread scheduler may assign tightly coupled threads to those processors which are in proximity to each other electronically, for example, processors in the same cluster, or processors in clusters that are in proximity to each other electronically if there are not enough available processors in one cluster.

8

compile the application into one or more object code which is specific to the underlying architecture. Additionally, the compiler analyzes data sharing behavior among different threads and obtains data sharing information before the threads are executed on target processors. The compiler then provides such data sharing information to the thread scheduler.

Thread scheduler 330 may be a part of, or associate with, an operating system (OS). Based on the data sharing information provided by the compiler, the thread scheduler may divide all of the threads into different groups. Threads within the same group share data with each other frequently. In one embodiment, the compiler may group all the threads into different groups based on their data sharing behavior obtained during compilation and pass the grouping information to the thread scheduler. Additionally, the thread scheduler is capable of obtaining structural characteristics of the underlying architecture of a multiprocessor system such as, for example, which processors are in the same cluster and which clusters are in proximity to each other electronically. If the OS has the knowledge of hardware structural features, the thread scheduler may obtain such hardware structural information from the OS. If the OS does not have direct knowledge of hardware structural features (e.g., in a situation where a virtual machine exists between the OS and the underlying hardware), the OS or the thread scheduler may still get such hardware information by invoking certain application program interfaces (API's).

With the hardware information, the thread scheduler may assign threads in the same group to processors in the same cluster. If there are not

enough processors available in any cluster, the thread schedule may identify a cluster that has the highest number of available processors, assign a corresponding number of threads in a group to processors in the cluster, and assign the rest of threads in the group to processors in another cluster that is in proximity to the cluster electronically. In case the cluster that has the highest number of available processors cannot host all of the threads in a group and does not have any cluster in proximity electronically that has enough available processors that can host the rest of the threads in the group, the thread scheduler may assign threads in the group to clusters that are in proximity to each other electronically and that together have enough processors available to host all of the threads in the group.

Feedback module 350 may observe actual data sharing behavior among threads during execution and provide this actual data sharing information to the thread scheduler. The thread scheduler may use such feedback information on actual data sharing behavior to regroup threads of the multithreaded application and reschedule the threads to target processors according to the regrouping result, at the next available thread scheduling time. The thread scheduler may first examine whether the feedback information supports regrouping and rescheduling. For example, if the actual data sharing behavior during execution does not deviate from the current grouping significantly, the thread scheduler may decide not to regroup or reschedule the threads. In one embodiment, the thread scheduler may decide whether to regroup and reschedule threads of the application based on the execution status of the application. For example, if the application is close

to completion when the feedback information regarding actual data sharing behavior is received and supports regrouping and rescheduling, the thread scheduler may choose not to regroup or reschedule threads at the next scheduling time.

Figure 4 is a flowchart illustrating an example process 400 for scheduling threads to target processors using information of data sharing behavior among different threads, according to the subject matter disclosed in the present application. At block 410, a multithreaded application may be received and passed to a compiler. At block 420, the multithreaded application may be compiled by the compiler. While compiling the application, the compiler may analyze data sharing behavior among threads of the application and collect data sharing information at block 430. Also at block 430, the data sharing information may be provided to a thread scheduler by the compiler. At block 440, threads that share data frequently, in other words, threads that are tightly coupled, may be identified based on the data sharing information provided by the compiler. All of the threads of the application may be placed into different groups, with those tightly coupled threads being in the same group. In one embodiment, the compiler may group all of the threads of the application using the data sharing information obtained during compilation and then pass the grouping information to the thread scheduler.

At block 450, threads in each group may be scheduled to processors that are in proximity to each other electronically. If there is any cluster that has enough available processors, threads in a group will be scheduled to processors in the same cluster; otherwise, threads will be scheduled to

processors in the minimum number of clusters that are in proximity to processors in the cluster. At block 460, threads that have been assigned may be executed on target processors. At block 470, a feedback module may observe actual data sharing behavior during execution on the target processors. At block 480 such actual data sharing information may be provided to the thread scheduler. Upon receiving such information, the thread scheduler will decide whether rescheduling of threads is necessary at the next scheduling time. If it is, the thread scheduler may regroup threads or fine tune previous grouping and reschedule the threads based on regrouping or fine tuning at the next available time.

Figure 5 illustrates an example program 500 that can be multithreaded in a multiprocessor system. This is a simple pointer chasing application. Depending on the size of the list p, it may take a long time to complete execution of this simple application. Thus, it may be necessary to multithread this application and run multiple threads in parallel. There are several ways to achieve parallel implementation of an application in multiple threads. One way is to use OpenMP, a parallel programming language, which has been supported by many compilers, such as Intel® C++ and Fortran compilers. The OpenMP programming model provides a work-queue execution model with task-queue and task primitives to allow users to efficiently exploit parallelism among irregular patterns of dynamic data structures and/or those with complicated control structures such as recursion. In the following description, the OpenMP programming model is used as an example to

illustrate how a data sharing aware thread scheduling scheme, as disclosed in this application, works.

Figure 6 illustrates a code 600 corresponding to the program shown in Figure 5, which is multithreaded for a multiprocessor system using the OpenMP programming model. The main differences between code 600 and code 500 are two added lines (Line 615 and line 635) in code 600, which parallelize the pointer chasing application shown in Figure 5.

Figure 7 illustrates an example assembly code 700 of the multithreaded program shown in Figure 6. Assembly code 700 is obtained by compiling code 600. Lines 750 through 770 are assembly translations for the master thread. Code 700 illustrates how thread scheduling is done without considering data sharing behavior among threads. Line 710 creates a team of threads. The compiler may use data sharing information among threads obtained during compilation to create a team of threads by placing those tightly coupled threads into the same team. However, assembly 700 does not provide a thread scheduler such data sharing information among threads. Thus, when the thread scheduler schedules a team of threads, it does not necessarily assign threads in the same team to processors in the same cluster. In other words, even if the data sharing information is considered in creating a team of threads, without informing the thread scheduler of such consideration, the resulting scheduling will not be based on the data sharing information obtained by the compiler.

Figure 8 illustrates an example assembly code 800 of the multithreaded program shown in Figure 6, which uses a data sharing aware thread

scheduling method, according to the subject matter disclosed in the present application. The main difference between code 800 and code 700 is line 815 which is not in code 700. In code 800, line 810 creates a team of threads with consideration of data sharing information among threads obtained by the compiler. Line 815 then informs the thread scheduler that threads in the same team should be scheduled to processors in the same cluster or in clusters that are in proximity to each other electronically.

Figure 9 illustrates an example code 900 of a multi-threaded video mining program, according to the subject matter disclosed in the present application. Code 900 is a sample from a video mining application. This application uses a task-queue to build a pipeline based parallel version. Within one task-queue, one thread is responsible for video decoding, and other worker threads perform feature extraction from decoded frames. In this application, the number of threads is obtained by invoking the function `omp_get_num_threads()`. Following the disclosed data sharing aware scheduling scheme, the compiler can easily add the scheduling hint information (e.g., line 815 in Figure 8) and the thread scheduler can then utilize this information to achieve efficient thread scheduling.

Figure 10 illustrates a performance improvement of a multithread application running on a multiprocessor system, obtained by using a data sharing aware thread scheduling module, according to the disclosed subject matter in the present application. An experiment was conducted by using a 32-way shared memory multiprocessor system to study the effectiveness of the disclosed data sharing aware thread scheduling scheme. Each processor

is equipped with an 8KB L1 cache, 512KB L2 cache, and 4M L3 Cache. Each cluster has 4 processors that share a 32MB L4 cache. The workload used for the study is a video mining application and a hybrid parallel scheme is employed, which exploits both data and task level parallelism to parallelize the application. With the hint information of data sharing provided by the compiler, the thread scheduler allocates threads in a team to processors in the same cluster (in this particular experiment, a team of threads has 4 threads).

For comparison purposes, two different cases are studied. In case 2, the disclosed data sharing thread scheduling scheme is used. In case 1, closely coupled threads, among which there is a lot of data sharing, are intentionally allocated to processors scattered in different clusters. Performance of each case is compared with the default scheduling scheme, which does not allocate threads without considering data sharing information among the threads. As shown in Figure 10, the x-axis shows the number of processors which are actually used in the multiple processor system; the y-axis shows the performance ratio between each of the two studied cases and the default scheduling scheme. Experiment results clearly show that performance in case 1 is consistently below performance in the default case for different number of processors. However, performance in case 2, which uses the disclosed data sharing aware thread scheduling scheme, is significantly better than performance in the default case for any number of processors studied in the experiment.

Although an example embodiment of the disclosed subject matter is described with reference to block and flow diagrams in Figures 1-10, persons of ordinary skill in the art will readily appreciate that many other methods of implementing the disclosed subject matter may alternatively be used. For example, the order of execution of the blocks in flow diagrams may be changed, and/or some of the blocks in block/flow diagrams described may be changed, eliminated, or combined.

In the preceding description, various aspects of the disclosed subject matter have been described. For purposes of explanation, specific numbers, systems and configurations were set forth in order to provide a thorough understanding of the subject matter. However, it is apparent to one skilled in the art having the benefit of this disclosure that the subject matter may be practiced without the specific details. In other instances, well-known features, components, or modules were omitted, simplified, combined, or split in order not to obscure the disclosed subject matter.

Various embodiments of the disclosed subject matter may be implemented in hardware, firmware, software, or combination thereof, and may be described by reference to or in conjunction with program code, such as instructions, functions, procedures, data structures, logic, application programs, design representations or formats for simulation, emulation, and fabrication of a design, which when accessed by a machine results in the machine performing tasks, defining abstract data types or low-level hardware contexts, or producing a result.

For simulations, program code may represent hardware using a hardware description language or another functional description language which essentially provides a model of how designed hardware is expected to perform. Program code may be assembly or machine language, or data that may be compiled and/or interpreted. Furthermore, it is common in the art to speak of software, in one form or another as taking an action or causing a result. Such expressions are merely a shorthand way of stating execution of program code by a processing system which causes a processor to perform an action or produce a result.

Program code may be stored in, for example, volatile and/or non-volatile memory, such as storage devices and/or an associated machine readable or machine accessible medium including solid-state memory, hard-drives, floppy-disks, optical storage, tapes, flash memory, memory sticks, digital video disks, digital versatile discs (DVDs), etc., as well as more exotic mediums such as machine-accessible biological state preserving storage. A machine readable medium may include any mechanism for storing, transmitting, or receiving information in a form readable by a machine, and the medium may include a tangible medium through which electrical, optical, acoustical or other form of propagated signals or carrier wave encoding the program code may pass, such as antennas, optical fibers, communications interfaces, etc. Program code may be transmitted in the form of packets, serial data, parallel data, propagated signals, etc., and may be used in a compressed or encrypted format.

Program code may be implemented in programs executing on programmable machines such as mobile or stationary computers, personal digital assistants, set top boxes, cellular telephones and pagers, and other electronic devices, each including a processor, volatile and/or non-volatile memory readable by the processor, at least one input device and/or one or more output devices. Program code may be applied to the data entered using the input device to perform the described embodiments and to generate output information. The output information may be applied to one or more output devices. One of ordinary skill in the art may appreciate that embodiments of the disclosed subject matter can be practiced with various computer system configurations, including multiprocessor or multiple-core processor systems, minicomputers, mainframe computers, as well as pervasive or miniature computers or processors that may be embedded into virtually any device. Embodiments of the disclosed subject matter can also be practiced in distributed computing environments where tasks may be performed by remote processing devices that are linked through a communications network.

Although operations may be described as a sequential process, some of the operations may in fact be performed in parallel, concurrently, and/or in a distributed environment, and with program code stored locally and/or remotely for access by single or multi-processor machines. In addition, in some embodiments the order of operations may be rearranged without departing from the spirit of the disclosed subject matter. Program code may be used by or in conjunction with embedded controllers.

While the disclosed subject matter has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments of the subject matter, which are apparent to persons skilled in the art to which the disclosed subject matter pertains are deemed to lie within the scope of the disclosed subject matter.

CLAIMS

What is claimed is:

1. A method for thread scheduling for a multithreaded application in a multiprocessor system, comprising:

obtaining information on data sharing behavior among multiple threads of said multithreaded application;

grouping said multiple threads into at least one group based at least in part on said information on data sharing behavior among said multiple threads; and

scheduling said group of threads to target processors, said target processors being in proximity to each other electronically.

2. The method of claim 1, wherein obtaining information on data sharing behavior comprises:

compiling said multithreaded application;

analyzing data sharing behavior among said multiple threads; and

collecting said information on data sharing behavior.

3. The method of claim 1, wherein grouping said multiple threads comprises:

identifying tightly coupled threads; and

placing said identified threads into the same group.

4. The method of claim 1, wherein scheduling said group of threads to target processors comprises:

assigning threads in the same group to said target processors that are in a cluster; and

if there are not enough available processors in said cluster, assigning the rest of the threads to processors that are electronically in proximity to said cluster.

5. The method of claim 1, further comprising:

executing said multiple threads of said multithreaded application on said multiprocessor system; and

observing data sharing behavior among said multiple threads during execution.

6. The method of claim 5, further comprising:

providing feedback of said data sharing behavior to regroup said multiple threads; and

rescheduling, at the next scheduling time, said regrouped multiple threads by assigning threads in the same group to target processors that are electronically in proximity to each other.

7. The method of claim 1, wherein grouping said multiple threads further comprises regrouping said multiple threads based on feedback of data sharing behavior among said multiple threads during execution.

8. An article comprising a machine-readable medium that contains instructions, which when executed by a processing platform, cause said processing platform to perform operations comprising:

obtaining information on data sharing behavior among multiple threads of said multithreaded application;

grouping said multiple threads into at least one group based at least in part on said information on data sharing behavior among said multiple threads; and

scheduling said group of threads to target processors, said target processors being in proximity to each other electronically.

9. The article of claim 8, wherein obtaining information on data sharing behavior comprises:

compiling said multithreaded application;

analyzing data sharing behavior among said multiple threads; and

collecting said information on data sharing behavior.

10. The article of claim 8, wherein grouping said multiple threads comprises:

identifying tightly coupled threads; and

placing said identified threads into the same group.

11. The article of claim 8, wherein scheduling said group of threads to target processors comprises:

assigning threads in the same group to said target processors that are in a cluster; and

if there are not enough available processors in said cluster, assigning the rest of the threads to processors that are electronically in proximity to said cluster.

12. The article of claim 8, wherein said operations further comprises:

executing said multiple threads of said multithreaded application on said multiprocessor system;

observing data sharing behavior among said multiple threads during execution;

providing feedback of said data sharing behavior to regroup said multiple threads; and

rescheduling, at the next scheduling time, said regrouped multiple threads by assigning threads in the same group to target processors that are electronically in proximity to each other.

13. The article of claim 8, wherein grouping said multiple threads further comprises regrouping said multiple threads based on feedback of data sharing behavior among said multiple threads during execution.

14. An apparatus for thread scheduling for a multithreaded application on a multiprocessor system, comprising:

a compiler to compile said multithreaded application, to analyze data sharing behavior among multiple threads of said multithreaded application, and to obtain information on data sharing behavior among said multiple threads; and

a thread scheduler to receive said information on data sharing behavior among said multiple threads and to schedule said multiple threads to target processors based at least in part on said information on data sharing behavior.

15. The apparatus of claim 14, wherein said thread scheduler groups tightly coupled threads into the same group and assigns said threads in the same group to processors in a cluster, and if there are not enough available processors in said cluster, assigns the rest of said threads to processors that are electronically in proximity to said cluster.

16. The apparatus of claim 14, further comprises a plurality of processors to receive and execute said scheduled threads.

17. The apparatus of claim 16, further comprises a feedback module to observe data sharing behavior among said scheduled threads during execution and to provide feedback on said data sharing behavior during execution to said thread scheduler.

18. The apparatus of claim 17, wherein said thread scheduler regroups and reschedules, at the next scheduling time, said multiple threads to target processors based on said feedback on said data sharing behavior during execution.

19. A multiprocessor system, comprising:

a synchronous dynamic random access memory ("SDRAM");

a plurality of processors coupled to access said SDRAM via a system interconnect, said plurality of processors forming at least one cluster, wherein processors in the same cluster have a shared storage device;

at least one compiler in one or more of said plurality of processors to receive and compile a multithreaded application, to analyze data sharing behavior among multiple threads of said multithreaded application, and to obtain information on data sharing behavior among said multiple threads; and

at least one thread scheduler in one or more of said plurality of processors to receive said information on data sharing behavior among said multiple threads from said compiler, and to schedule said multiple threads to said plurality of processors based at least in part on said information on data sharing behavior.

20. The system of claim 19, wherein said thread scheduler groups tightly coupled threads into the same group and assigns said threads in the same group to processors in a cluster, and if there are not enough available

processors in said cluster, assigns the rest of said threads to processors that are electronically in proximity to said cluster.

21. The system of claim 19, further comprises a feedback module to observe data sharing behavior among said scheduled threads during execution and to provide feedback on said data sharing behavior during execution to said thread scheduler.

22. The system of claim 21, wherein said thread scheduler regroups and reschedules, at the next scheduling time, said multiple threads to said plurality of processors based on said feedback on said data sharing behavior during execution.

ABSTRACT OF THE DISCLOSURE

According to embodiments of the subject matter disclosed in this application, a compiler in a multiprocessor system may compile a received multithreaded application, analyze data sharing behavior among multiple threads of the multiprocessor application, and provide such information to a thread schedule in the multiprocessor system. Threads that share data frequently may be grouped together based on the data sharing information provided by the compiler, and at run time, the thread scheduler may schedule threads in the same group to processors that are in proximity to each other electronically. Additionally, a feedback module may collect information on data sharing behavior among threads during execution and feedback such information to the thread scheduler. The thread scheduler may use the feedback information to regroup and reschedule the threads at the next available scheduling time.





FIGURE 1



300




FIGURE 3

4/8




FIGURE 4

500

```
510 void test(List p)
520 {
530     while (p != NULL)
540     {
550         do_work(p);
560         p = p->next;
570     }
580 }
```

FIGURE 5

600

```
605 void test(List p)
610 {
615     #pragma intel omp parallel taskq shared(p) num_threads(4)
620     {
625         while (p != NULL)
630         {
635             #pragma intel omp task captureprivate(p)
640             {
645                 do_work(p);
650             }
655             p = p->next;
660         }
665     }
670 }
```

FIGURE 6

700

```

705 __kmpc_serialized_parallel;
710 __kmpc_fork_call ; create a team of threads
715 __kmpc_taskq ; select one thread to be master thread
720 __kmpc_end_taskq ; turn others to worker threads
725 __kmpc_end_serialized_parallel

750 __kmpc_task_buffer ; allocate memory for thunk_t
755 copy captureprivate variable to thunk_t
760 __kmpc_task ; enqueueing of a task
765 __kmpc_taskq_task ; enqueueing the task_thunk
770 __kmpc_end_taskq ; turn into worker thread

```

FIGURE 7

800

```

805 __kmpc_serialized_parallel;
810 __kmpc_fork_call ; create a team of threads
815 __kmpc_team_scheduling ; schedule team of threads on nearby processors
820 __kmpc_taskq ; select one thread to be master thread
825 __kmpc_end_taskq ; turn others to worker threads
830 __kmpc_end_serialized_parallel

850 __kmpc_task_buffer ; allocate memory for thunk_t
855 copy captureprivate variable to thunk_t
860 __kmpc_task ; enqueueing of a task
865 __kmpc_taskq_task ; enqueueing the task_thunk
870 __kmpc_end_taskq ; turn into worker thread

```

FIGURE 8

900

```
905 int NumThreads; // specify the number of threads
910 int NumThreadsPerGroup; //specify the number of threads per group
915 int GroupNum = NumThreads / NumThreadsPerGroup;
920 #pragma omp parallel for num_threads(GroupNum)
925 {
930     for (int gid = 0; gid < GroupNum; gid++)
935     {
940         #pragma Intel omp parallel taskq num_threads(NumThreadsPerGroup) ordered
           shared(frame)
           {
945             #pragma Intel omp ordered
           frame = decoding_video();
950             #pragma Intel omp task captureprivate(frame)
           extract_feature(frame);
955             }
960         }
965     }
970 }
975 }
980 }
```

FIGURE 9




FIGURE 10

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☒ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.